

-
- Current Meeting Report
 - Slides
-



2.3.10 Networks in the Small - aka Home Networks (nits) bof

Current Meeting Report

Minutes: NITS BOF

15 March, 1999

Minneapolis IETF

CoChairs:

Stuart Cheshire, cheshire@apple.com

Peter Ford, peterf@microsoft.com

Mailing list: nits@merit.edu

After a short welcome and discussion of agenda we established a short list of presentations to discuss home networks. It appears that about 180-200 people attended the BOF, with a full room.

The list of presentations included:

- 1) Overall Goals of NITS BOF - Peter Ford
- 2) Home Networking Device and Service Discovery Requirements
- draft-miller-homedisc-req-00.txt - Brent Miller
- 3) Automatic IP address assignment for link local address w/IPv4
- draft-ietf-dhc-ipv4-autoconfig-03.txt - Stuart Cheshire/Ryan Troll
- 4) Multicast Discovery of DNS Services
- draft-manning-multicast-dns-01.txt - Bill Woodcock/Bill Manning
- 5) Service Location Protocol
- draft-ietf-svrl-protocol-v2-12.txt - Erik Guttman
- 6) Simple Service Discovery Protocol
- draft-cai-ssdp-v1-00.txt - Yaron Goland
- 7) IPv6 Autoconfiguration and Neighbor Discovery - T. Narten
- 8) IP Address sharing - Stuart Cheshire

Goals of NITS BOF

Presented by Peter Ford and Stuart Cheshire

Goals articulated by Peter:

- 1) Discuss How Networks in the Small (e.g. Home Networks) can and should work
- 2) Determine if there is standards work to be done
- 3) If so, charter appropriate work
- 4) Non-goal - solve all problems by 5.30pm

Why should the IETF consider working on NITS now?

- 1) IP stack has sufficiently standardized and many vendors are looking at IP based devices that should have "plug and play behaviour".
- 2) IP currently fails the simplicity test - if one were to go to a CostCO and buy 2 PCs and an ethernet hub someone would still have to configure the IP addressing and other stack parameters.

Peter suggested that we should benchmark current IPv4 practices against those of Appletalk, IPX, and NetBIOS, including deployment of DHCP servers and bootstrapping thru DHCP to get DNS server addresses and domains configured in an end system. When one considers that to print many people then have to go edit their /etc/printcap and /etc/resolv.conf files we should turn red in embarrassment. With Appletalk and NetBIOS over NetBeui or IPX you can bring a system up and print with minimal (close to zero) configuration.

Beyond simple addressing and DNS configuration we also need to worry about router, proxy and NAT configs that get more daunting every day.

Brent Miller from IBM Raleigh went through a brief description of the requirements his team from IBM developed for Home networks.

Basic assumptions include that there is a home LAN that is intermittently connected to the Internet and the LAN and LAN clients use standard IETF protocols. A computer that is introduced to this LAN should be able to enjoy basic connectivity to other similar systems and to services on that network. IN short, things should "just work".

Brent discussed a taxonomy of requirements from the draft including:

Autoconfiguration: IP address assignment to new devices, Service/Device location, and the use of User friendly names.

Question: Is this limited to single subnet single domain? What about multiple administrative domains? What if your teenage children want to run their own autonomous domains at home?

Answer: In order to make progress, should not require NITS solutions to scale beyond single subnet, single administrative domain.

Automatic IP address assignment for link local address w/IPv4
- draft-ietf-dhc-ipv4-autoconfig-03.txt - Stuart Cheshire/Ryan Troll

Stuart Cheshire presented a basic overview of IPv4 address self-configuration as currently implemented by Apple and MS. The purpose is to support configuration of basic stacks without manual configuration. Stuart pointed out that both Apple and Microsoft are attempting to move towards complete use of IETF standard protocols in replacement of their legacy protocols (Appletalk and NetBeui).

Stuart started off his slide deck with a simple description of autonet:

- 1) Designed for small isolated LANs
 - a home LAN
 - a small office
- 2) Currently in Windows 98 and MAC OS 8.5
- 3) Described in an Internet-Draft by Ryan Troll (Carnegie Mellon).

Stuart proceeded to describe the operation in Mac OS 8.5.

- 1) DHCP Discover
- 2) If no reply retry, 4, 8, 16 second intervals
- 4) If no DHCP server discovered then?

4a) Pick a random address on 169.254.*.* subnet (except first and last 256 addresses which are reserved for future use).

Send an ARP probe (ala DHCP conflict detection) to verify address is not already in use

If address in use, go back to 4a) - iterate at most 10 times and then fail stack initialization

Configure the interface with IP address, and start using interface

Every 5 minutes send single DHCP Discover to determine if a DHCP server has come online on the LAN, if so then proceed to normal DHCP client behavior upon DHCP Offer message reception.

Stuart then pointed out key points/features:

- 1) allows invisible use of IP - a user can go to the chooser using appletalk, but actually connects to it using autoconfigured TCP/IP
- 2) This is NOT a substitute with IPv6 - IPv6 is critical so that everyone can get globally unique public IP addresses
- 3) Stuart would like to also perform resource and service discovery using IP instead of Appletalk Name Binding Protocol (NBP).

There were questions about some of the DHCP specifics. Also several questions came up on who "owned" the 169.254/16 subnet, and Bill Manning offered that the IANA did at that is what WHOIS tells people who see this happening behind their firewalls!

MUDDS

Multicast Discovery of DNS Services

Presented by Bill Woodcock

Bill indicated that version 1 of the spec is in the Internet Draft directory but that he and Bill Manning are up to version 3 of the spec. He noted that people are looking for simple replacements for existing protocols such as Appletalk NBP and NetBIOS browsing.

There are several applications:

- 1) Bootstrapping DNS resolvers on clients when there is not a DHCP server
- 2) Discovery of unconfigured devices such as printers and routers
- 3) Let Apple deprecate the AppleTalk Chooser

Bill also noted that this mechanism could be used to provide a lightweight subset of SLP features. Steve Deering asked how this works and Bill briefly described how you can find SRV records on hosts that are listening to the multicast port. A question was asked about the hostname part of the SRV record if you were looking for "any printer on the LAN" and Bill noted that DNS supports the use of wildcards (*) in the names, something which raised several eyebrows in the audience. Erik Guttman asked about how do you get DNS response suppression by end nodes if a DNS server WAS present on a LAN to which Bill described how one would first look for a DNS server via multicast and if a server was found you would only use unicast queries to that server. This raised several questions about the scoping of multicast requests and several references to the work in the multicast working groups on administratively scoped multicast.

Erik Guttman noted that you need to put in aggressive backoffs into the protocol to prevent swamping of the LAN, especially in wireless cases. He also noted that there were issues of trust models, which are contained in SLP.

Henry Sinnreich of MCI asked several questions about whether this mechanism could be extended to find public DNS servers. This was answered that it depends on careful configuration of multicast scopes.

Bill pointed out that the major advantage of using multicast for DNS discovery is that it uses pre-existing technologies: Multicast and DNS. The multicast usage employs a statically assigned address within the administrative local-scope and SRV records are used to describe network services such as DNS or printing. The DNS transaction works as normal except the initial query is performed within the multicast group rather than with a preconfigured DNS server.

SLP Overview by Erik Guttman

Erik gave a brief overview of how SLP works and the current standards status. He mentioned several small last minute additions to cover the cases where people are using directories such as LDAP and what a minimal subset of SLP would be. He also described some future work in the SLP group that would address issues that some parties have had with the use of SLP in environments with LDAP servers:

- 1) an LDAP server could emit DA Adverts allowing LDAP clients to use LDAP directly instead of using SLP front ending a DS.
- 2) Erik described how JINI and non-JINI services are discovered using SLP. Erik then compared SLP with multicast DNS:

Queries in SLP are of the form: "servers that ..." such as what are the printers that all have 721 dpi support.

SLP is carefully crafted to work in networks from small to large

Steve Deering proceeded to ask if any consideration was made in SLP to take advantage of many multicast addresses so that the multicast IP filter could filter out many questions irrelevant to the server instead of waking all servers with all SLP queries. Erik said this was in the early versions of the protocol, but that it was pulled out of SLPv2 due to no apparent interest.

Charlie Perkins noted that SLP was designed against a set of requirements that looked surprisingly similar to the requirements that Brent Miller presented at the beginning of the session.

Simple Service Discovery presented by Yaron Goland

Yaron discussed Microsoft's investigations into the subset of home networking including service discovery.

The problem statement is how systems can discover each other even if there is no network administrator or directory support. A subset of this problem is to discover directory support if present on the network.

The target environment anticipated is:

- 1) ~~an IP network with multicast support~~
 - 2) limited memory and storage - aggregate of 1M of memory or less
 - 3) ~~HTTP and XML support expected to be common for supporting device to device communication.~~
- Yaron noted that investigation of embedded web systems web sites indicate that web servers can be small (< 70Kbytes of code).
- 4) Example devices include: thermostats, security systems, CD players, printers, scanners, etc.

Solution Parameters include:

- 1) ~~multicast IP is used to enable discovery in the absence of directory.~~
- 2) HTTP/XML based - to re-use stacks already present in the devices (small web servers that support UI, mgmt, etc.)
- 3) no parameters - services are identified with URIs, any more powerful discovery or parameter negotiation is done in a "higher" protocol or a service specific protocol (e.g. IPP).

Yaron noted the draft on SSDP in the Interdraft directory and also noted there were several items TBD. A proposed implementation is:

- 1) ~~use HTTP over multicast UDP~~
- 2) Declaration based discovery using OPTIONS method and If header
- 3) ~~announcement based discovery using ANNOUNCE method and resource-state header. This would allow new services to introduce themselves on the fly.~~

There was good followup discussion on the main differences btw SSDP and SLP. SSDP does not support queries of the form "who are the XXX servers that are ...", only dealing with "who are XXX servers?"

IPv6 autoconfiguration and link local addressing - Tom Narten

The goal is to present an overview of what was done and standardized in IPv6 to support small networks

For autoconfiguration IPv6 provides for:

PNP out of the box experience including support for the isolated dentist's office (single LAN that is not Internet connected) with support that scales to large enterprise.

All addresses are "leased"

IPv6 has built in support for multiple addresses per interface.

IPv6 also has scoped addressing:

- 1) Link-Local for communication with immediate neighbor
- 2) Site-Local which is analogous to IPv4 net 10/8
- 3) global addresses that are globally unique

These are all documented in RFCs (question from audience).

Link Local addresses are such that every node assigns link local addresses to its interfaces, those addresses are only good for that link/subnetwork. Link local addresses have a well known prefix (fe80::/9). Each address has an interface identifier that is derived from the MAC address, which is globally unique. There is built in duplicate address detection. It works much like the IPv4 autoconfiguration with a much lower probability of collision - probably none.

IPv6 supports both DHCP and stateless address autoconfiguration where there is no server on the wire (usually the router). A router sends a router advertisement and the end system cons's up an address from a prefix contained in the router advertisement (RA) and the Interface ID derived from the MAC of the Interface. If the interfaces sees multiple RAs then it means the host can have multiple addresses. One changes an interfaces address by changing the prefix contained in the RA. Tom noted that IPv6 uses dynamic DNS to update the DNS when an interface's address changes.

Router advertisements contain a default router list. Each host picks this up and maintains a local list of default routers. The host keeps track of reachability to all neighbors, if necessary the host can send probes.

IP address sharing

Stuart Cheshire/Apple Computer, Inc.

Stuart began describing the current status of getting PacBell ADSL at up to T-1 data rates. To get a single IP address for a single host the service is \$50/month. If you want more IP hosts then it costs >\$100 per month.

Stuart infers that IPv4 addresses are becoming scarce.

People are finding an arbitrage around additional cost per address in the deployment of NATs. He notes this is not an end-all solution. Nats are: fragile, break the end to end model, breaks IPSEC, requires separate support on a per protocol basis such as ftp and any other protocol that buries ports inside PDUs. The result is that NATs hold state on a per connection basis and is a single point of failure.

How could you make hosts behind NATs work better if they did know they were behind a NAT? Stuart poses that there can be address sharing aware (AAA) hosts. You need to ensure that on a single LAN where you are sharing the same IP address across hosts that hosts use unique to the LAN ports. The

NAT would need to be able to demultiplex inbound packets to hosts based on which hosts on the LAN have what ports in use. To that end Stuart proposed an "extended ARP" that is based on <IP address, proto, port> instead of simply using the IP addr. This can be used as a claim mechanism as well as used by the NAT box to deliver a packet to the right host.

Stuart wanted to make it clear this was a 1-2 year solution and that IPv6 should be the real answer in the long term when it is fully deployed.

Several members in the group noted similar ideas were being discussed in the NAT working group and perhaps this topic should be out of scope of NITS due to active work in the NAT WG.

Closing Discussion

There appears to be sufficient interest in the area of small networks to proceed with forming a work list. Stuart and Peter took the action to work out a charter with the ADs (Thomas Narten and Erik Nordmark).

In enumerating the topics to be discussed we collected:

- 1) multicast DNS: more spec work needed, name resolution for IPv4 and v6, service discovery. Many voiced concern over m'cast scoping of queries. This topic had a lg amount of support for wg activity.
- 2) service discovery? mcast DNS for this purpose vs. SLP or SSDP? Issues to be determined was the spectrum of scalability, would mcast DNS suffice? Some mentioned applicability as an issue to be resolved.
- 3) What about IPv6, perhaps we should skip v4 and move directly to v6.
- 4) Need to refine requirements doc to aid in resolution of item 2) above.
- 5) What about multiple LANs (e.g. wireless, 1394 and ethernet all in one house)? Most felt we needed to worry about this topic.
- 6) what about multiple administrative domains on the same LAN such as in apartment buildings?
- 7) security was absent for most of the presentations, what are the security models and there was some who shared that they believed that security in the home was VERY important.
- 8) Is mobility an issue?
- 9) what about intranet vs Internet as a connectivity model (some said the use of the word intranet should be grounds for banishment ...). This appeared in the context of should the addresses on Home LANs be private or public, and how should 2 homes interconnect if private addresses are used

The meeting drew to a close at 6pm.

Slides

None received.